

REMARKS/ARGUMENTS

Claims 1-20 are pending. Claims 8-10, 12-16, and 18-20 have been amended to fix typographical errors in the reference to their respective independent claims. No claims have been canceled.

35 U.S.C. § 102 Rejection, Wellen.

Claims 1-20 were rejected under 35 U.S.C. § 102(e) as being anticipated by Wellen (US patent publication 2002/0075884). Reconsideration and allowance of claims are respectfully requested for the following reasons.

Claims 1-6

Wellen does not anticipate claim 1 because Wellen does not teach or suggest each and every element of claim 1. For example, claim 1 recites, among other elements, "identifying a first portion of a first message in a first slice of a switch", "identifying a second portion of the first message in a second slice of the switch.", "selecting, independently in each slice" a message, and "sending the first portion of the selected message from the first slice to the network resource specified by the ... routing portions corresponding to the selected message; and sending the second portion of the selected message from the second slice to the network resource specified by the ... routing portions corresponding to the selected message."

Nowhere does Wellen show or suggest each and every of these recited features. In particular, Wellen's description in paragraphs 29-31 of (1) a switch that connects input data streams to output data streams and of (2) a virtual queue working in conjunction with a scheduling function do not describe a slice of a switch as recited in the invention, do not describe each slice selecting a portion of message independently, and do not describe a first slice and a second slice handling different portions of the same message, which are then both sent to the same network resource.

Present Invention

The present invention overcomes the limitations on the size of a typical crossbar switch and the restrictions on the length of a data word that can be switched by conventional crossbar switches, thus providing good scalability with a large number of processors. See page 1, lines 16-20 of the specification. To accomplish this, first the processor crossbar PXB divides a message into multiple portions so that the size of each portion is not too large for a switch crossbar SXB (page 6, lines 25-27) and to allow for parallel transfer. Second, a data bus simultaneously transfers portions of a message to different slices of a switch, which will then identify the portions. See page 10, lines 15-19. In order to store and select different portions of different messages, each slice includes a buffer and a multiplexer, as well as an arbiter, so that the selection is done "*independently in each slice*," as recited in claim 1. See page 6, lines 7-8. This step prevents bottlenecks in the transfer of messages to the switch, if, for example, one processor group was sending many more messages to the same switch crossbar than other processor groups.

Third, within a switch and memory crossbar, the movement of different portions of a message is handled by different slices. See Figures 3 and 4. In other words, portions of the message are transferred in parallel through the switch crossbar. See page 10, lines 15-19 and page 3, lines 5-7. The different portions of a message are transferred to the same resource. See page 6, lines 25-30. Claim 1 recites "*sending the first portion ... to the network resource specified by the... routing portions corresponding to the selected message; and sending the second portion ... to the network resource specified by ... routing portions corresponding to the selected message*." This step enables more efficient throughput due to the parallel operation of the slices on different parts of the same message. In transferring the portions of the message, a first slice identifies a first portion of the message and the second slice identifies a second portion of the message, which is a natural result of the parallel movement of the different portions of a message through the switch crossbar.

No Slices that Each have Storage and Selection Capability

Wellen does not teach or disclose separate components, such as a slice, where each includes a buffer, a multiplexer, and an arbiter. Therefore, there is no disclosure of a step involving such a slice, particularly a step of each slice independently selecting a portion of a message.

Paragraph 29 discusses switching input data streams to specific output data streams by directly connecting input ports to the scheduled output ports. There is no discussion of storing data packets of a data stream, which is the function of the buffer in a slice. Since there is only a choice on the scale of which input data stream is connected which output data stream, there is also no discussion of selecting which packet may be transferred to an output port as does the arbiter. Thus, paragraph 29 does not teach or disclose a slice of a switch.

Paragraph 31 does discuss virtual queues and a scheduling function, which has an embodiment in Figure 2. However, the separate queues 6 in figure 2 along with the scheduler 7 cannot each be a slice because Wellen does not have a separate scheduler for each queue. In the present invention, there is a separate arbiter for each buffer. Thus, each queue in Figure 2 does not teach or disclose a slice of a switch, particularly slices that select a portion of a message independently in each slice.

Note that the case of the entire structure of Figure 2 being interpreted as a slice does not allow for different portions of the same message to be identified by different slices, which is discussed in the next section with reference to paragraphs 30 and 31.

No Identifying of Different Portions of the Same Message in Different Slices

Wellen does not teach or suggest a first portion of a first message in a first slice of a switch identifying a second portion of the first message in a second slice of the switch.

At paragraph 29, Wellen describes "an application of a scheduler for switching input data streams to specific output data streams. A switch 2 connects N_i input ports . . . to N_o output ports" A "stream" is anything issuing or moving with continued succession of parts must be used. In this description, the data packets, in order to succeed each other, must follow each other serially from N_i to N_o . The identification of two portions (data packets) of the same

message in different slices cannot occur due to this serial movement. Thus, there is no step of a first slice identifying a first portion of the message and a second slice identifying a second portion of the message.

If each virtual queue of paragraphs 30 and 31 (or group of queues with a single scheduler as in Figure 2) for each output/input port combination are interpreted as a slice, then different portions of the same message cannot be identified by different slices and then sent to the same network resource. Wellen describes putting data packets of an input data stream bound for a specific output port into one queue (slice). Note that to be portions of a message, they are sent to the same network resource as in claim 1. Because of this and Wellen's construction of the queues, the data packets from the same message must be bound for the same output port. Thus, in Wellen, all of the data packets from the same message transfer through the same queue (slice), which prohibits identifying different portions of the same message in different slices.

For these reasons, the cited reference does not show or suggest the invention as recited. Accordingly, claim 1 and its dependent claims are allowable.

Claims 7-10

Claim 7 recites similar features as recited for claim 1, and claim 7 and its dependent claims should be allowable for at least similar reasons as claim 1.

Claims 11-16

Claim 11 recites similar features as recited for claim 1, and claim 11 and its dependent claims should be allowable for at least similar reasons as claim 1.

Claims 17-20

Claim 17 recites similar features as recited for claim 1, and claim 17 and its dependent claims should be allowable for at least similar reasons as claim 1.

35 U.S.C. § 102 Rejection, Chiussi et al.

Claims 1-20 were rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 5,689,506 to Chiussi et al. (Chiussi). Reconsideration and allowance of claims are respectfully requested for the following reasons.

Claims 1-6

Chiussi does not anticipate claim 1 because Chiussi does not disclose or suggest each and every element of claim 1. For example, claim 1 recites, among other elements, *"identifying a first portion of a first message in a first slice of a switch", "identifying a second portion of the first message in a second slice of the switch.", "selecting, independently in each slice" a message, and "sending the first portion of the selected message from the first slice to the network resource specified by the... routing portions corresponding to the selected message; and sending the second portion of the selected message from the second slice to the network resource specified by the ... routing portions corresponding to the selected message."*

Nowhere does Chiussi show or suggest these recited features. In particular, Chiussi's description of replicating a cell and transferring it to multiple output ports does not describe a slice of a switch as recited in the invention, does not describe each slice selecting a portion of message independently, and does not describe a first slice and a second slice handling different portions of the same message, which are then both sent to the same network resource.

No Slices that Each have Selection Capability

Chiussi does not teach or disclose separate components, such as a slice, where each includes a buffer, a multiplexer, and an arbiter. Therefore, there is no disclosure of a step involving such a slice, particularly a step of each slice independently selecting a portion of a message. At column 9, lines 56-60, Chiussi describes transferring data cells from a queue to the switch fabric. There is selection logic that works on a round-robin basis for determining the next cell to transfer from all of the queues. The queues are all serviced by the same selection logic, which prohibits each queue being a slice since each queue does not have its own arbiter. Claim 1 provides the selection is done *"independently in each slice,"* and Chiussi does not have this. Additionally, a single selection logic services the transfer of cells from the switch fabric to the output ports, thus these cannot be different slices either. See column 12, lines 16-23. Thus, Chiussi does not disclose each slice independently selecting a portion of a message.

No Identifying of Different Portions of the Same Message in Different Slices

Chiussi does not teach or suggest a first portion of a first message in a first slice of a switch identifying a second portion of the first message in a second slice of the switch.

In column 8, lines 52-54, Chiussi states: "The delivery of the multicast cells occurs through a virtual connection (VC) established between input port 102 and the requested output ports." The specific use of port 102 is exemplary, but the illustration does provide that different cells of a multicast connection come into the switch through the same input port. There is no discussion of different cells of a multicast connection being identified in different input ports (slices). In other words, there is no parallel transfer of the cells from the same multicast session through the input ports. Since Chiussi discusses using the same translation table of a multicast request to define the path of a cell through the switch (column 4, lines 1-18 and column 8, lines 54-61), further stages of the switch do not disclose or suggest cells (portions) from the same multicast session (message) bound for the same network resource (output port) being identified in different stages (slices). Thus, Chiussi does not disclose or suggest identifying a first portion of a first message in a first slice, identifying a second portion of the first message in a second slice, and sending the portions to the same network resource.

Accordingly, claim 1 and its dependent claims are allowable.

Claims 7-10

Claim 7 recites similar features as recited for claim 1, and claim 7 and its dependent claims should be allowable for at least similar reasons as claim 1.

Claims 11-16

Claim 11 recites similar features as recited for claim 1, and claim 11 and its dependent claims should be allowable for at least similar reasons as claim 1.

Claims 17-20

Claim 17 recites similar features as recited for claim 1, and claim 17 and its dependent claims should be allowable for at least similar reasons as claim 1.

Appl. No. 09/925,156
Amdt. dated January 5, 2005
Amendment under 37 CFR 1.116 Expedited Procedure
Examining Group 2112


PATENT

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400, extension 5213.

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'Melvin D. Chan', with a long horizontal flourish extending to the right.

Melvin D. Chan
Reg. No. 39,626

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 650-326-2422
MDC:dbk/km
60379488 v1